




THINKING INSIDE THE BOX

ILLUSTRATION: MILES COLE

Computers that think like people aren't just the stuff of science fiction: they are being developed right now. But to simulate us properly they need to grasp human experiences. Jon Thompson explains how machines could get to know us better

Hollywood rarely paints a rosy picture of artificial intelligence – 2001, *Terminator* – it seems every time we give computers the ability to think for themselves, they decide they're better off without us. Yet artificial intelligence, or AI, is already with us. Speech recognition software has to exercise some intelligence to understand different dialects and differentiate between "a typical" and "atypical", for example. AI is also employed in fraud detection, trend analysis, face and optical character recognition, and many other areas.

So-called weak AI systems, which can replace humans in specific, repetitive tasks, are relatively easy to design. Strong AI, however, seeks to make machines react intelligently to the world in general, and this isn't so straightforward. Strong AI requires telling software a lot about the world first. You pull strings, but you can pull or push a lever. Water always flows downhill because of gravity. Night follows day.

All these things are common sense to us, but a computer doesn't have real-life experiences from which to create an understanding of the world. Instead, it needs to be fed with vast amounts of information. Done manually, this would take decades to implement in sufficient detail to be useful. Luckily, serious efforts began over 20 years ago.



ensure they captured this more subtle information, they introduced the idea of transivity. Suppose that page A linked to page B, and page B linked to page C. Pages A and C obviously have a relationship to each other, but it is weaker than the relationship between A and B, and B and C. Transivity allows the web server to record a new, weaker link directly between A and C.

The server also used a rule to enforce the symmetry of the strengths of the links. So if page A related strongly to page B, the inverse must also be true. During the experiment, symmetry worked on links created by the transivity rule too.

When the experiment ended, Bollen and Heylighen found some interesting patterns in the strengths of the connections between pages. After only 150 clicks, the network began to show clear signs of self-organisation into a knowledge base. Soon it was possible to see how related topics cluster together naturally. Pages linked to most frequently from the 'time' page included age, century, day, evening, moment, period, week and year. Those visited most frequently from the 'space' page included place, area, point and stage.

The resulting paper (<http://pespmc1.vub.ac.be/papers/selforganwww.html>) puts forward some fascinating ideas for applying this self-organising principle to the World Wide Web. Each request for a web page that a user makes carries a referrer string, which allows the target page to see where the visitor is browsing from; this takes care of symmetry. If a future version of the HTTP protocol were to include a previous referrer string, transivity could be analysed across the gigantic mesh of web servers the world over, as well as between the pages they contain.

AN OPEN MIND

One project that has benefited directly from the web is Open Mind Common Sense (OMCS), based at MIT's Media Lab. Rather than attempt to process existing web pages, OMCS draws on the vast numbers of web users who are willing to contribute to a worthy scientific cause.

While Cyc's knowledge base was fed by many contributors, its CycL language had to be learned before contributions could be added. A natural language interface would mean that anyone could join in. Weak AI is good enough at natural language to extract meaning from specially formatted sentences that it can put into a knowledge base. There's a web interface to OMCS so anyone can register and tell it new things about the world. The website contains random examples of things people have told it, such as, for example, that motorcycles have two wheels and an engine, but other bikes have pedals. Anyone can join in by visiting <http://commonsense.media.mit.edu>. OMCS contains over 700,000 pieces of information, which have found uses in a number of other projects, including the Roboverse Domain project.

The Roboverse uses the OMCS knowledge base to give common sense to a group of virtual, mannequin-style robots living in their own computer-generated world. Users give the robots tasks to complete by typing in plain English. The robots engage in reasoning about the position and physical properties of objects. Furthermore, they also call on each other for help in completing their goals.

If, for instance, one robot is building a tower from blocks and cannot reach something he

Semantic gesture

When Sir Tim Berners-Lee gets an idea, the world tends to listen. After all, his last one was the World Wide Web. His thoughts now involve turning the web into an artificially intelligent system he calls the Semantic Web.


The power contained within the Semantic Web lies in a deceptively simple idea. Web page owners add a series of special XML tags called a Resource Definition Framework (RDF) to their web pages. These RDF tags point to entries in machine-readable databases of concepts. By referencing these unique definitions, a semantic search engine can search the web for other instances of the same reference. In effect, it searches by meaning rather than keyword. It can widen its search by looking for concepts indicated as related to the original in the online definition database, too.

The definitions used by the web page tags are unambiguous and can be highly detailed.

The idea is that software will be able to understand how to use the facilities provided by a website from its RDF tags. This could lead to programs that reason about what they find and can do everything from research to our weekly shopping. It sounds great, but so far RDF has attracted few developers.

Since its announcement in 2002, only 200,000 websites out of around eight billion have become Semantic Web-enabled. That's about 0.005 per cent of the total. Also, many page authors have abandoned their work over the years, and long forgotten their web space passwords, so it's going to be difficult to add RDF to a large proportion of the web.

There are also difficulties in deciding which are the right tags to add to a particular web page. After all, in describing a controversial news story, one man's terrorist is another man's freedom fighter.



The great thing about MIT's Open Minds Common Sense knowledge base is that anyone can join in and add their own observations about the world

wishes to use, he can ask another robot for help. The second robot can work out what help to give. The user can direct them all to perform new tasks and can ask why they are doing certain things. The Roboverse also models physical forces such as gravity, adding more complexity to tasks. If one robot asks another that is already engaged in activity for help, the second robot will have to reason about conflicts of interest and opportunities for cooperation.

GOOD FOR THE GOOSE

Another MIT OMCS project is GOOSE, a common sense-based search engine by Hugo Liu, Henry Lieberman and Ted Selker. Although engines such as Google and Yahoo! do an excellent job of indexing billions of web pages, they only match the keywords you enter with those found in their databases. Their technology is smart and fast, but there's only a limited capability for thinking about what your search actually means.

Search engines tend to use one of three techniques to add intelligence to their searches. Some use a thesaurus to try words that mean nearly the same as the original keyword. Others collect data on the relevance for the user of the

pages they return, so they can associate pages and keywords better. They do this by inviting those who use the search engine to score their results. Some, promoted as having a natural language interface, look for question forms they recognise and from which they can extract keywords.

GOOSE goes further, however, applying its understanding of the world to arrive at a deeper understanding of the request. Typing "Help me get rid of the mice

from my kitchen" will cause it to reason about mice, the various contexts in which they exist, and what it means when they appear in a kitchen. Correctly realising that they must be pests and not pets, it searches for pest control websites. In the paper *GOOSE: A Goal-Oriented Search Engine with Commonsense* (<http://web.media.mit.edu/~lieber/publications/goose.pdf>), its developers say GOOSE is able to do this because of all the common sense it carries. From this, it generates its own search query.

The paper goes on to describe possible future versions of the software, where the request "Find a birthday present for my brother" might return pages selling furniture. The program would do this because it knows that your brother has just left home, people who have just left home need to buy furniture, furniture is expensive, and buying an expensive gift shows you care.

OMCS and the Roboverse are impressive, but where exactly does this leave the original Cyc project? Is it lagging behind for want of a natural language interface to speed the flow of new information? Far from it, in fact. Cyc is now learning to gather facts from the web using natural language.



BANANA BRAIN

If we were to draw a large diagram showing how different objects, actions and concepts relate to each other, we could encode a lot of information about the world. By using many different link types, we could describe what we knew about the world in fine detail. And if we were to create this diagram in an electronic form, it would be relatively easy for a computer to understand.

Concepts that are several hops away on our diagram are less relevant than those connected directly; we can also substitute well-trodden multi-hop paths with direct links using transitivity. In other words, we can have our diagram discover new associations. All we need then is a large source of common sense from which our computer can learn automatically.

In a paper published in November 2004, Rudi Cilibrasi and Paul Vitanyi from the University of Amsterdam suggest that we can automatically generate a knowledge base so richly interwoven that it captures how all the concepts it contains relate to the others, no matter how large it is (see <http://www.arxiv.org/abs/%20cs.cl/0412098>). All that's needed, they say, is the public programming interface to Google.

The underlying idea is remarkably simple. It uses only the number of hits produced by a search for pages containing pairs of words. Searching for the words 'horse' and 'rider' produces 2.9 million hits, but searching for 'horse' and 'banana' produces just 750,000. From this, we can see that horses have far more to do with riders (or oats, saddles, racing or gambling) than they do bananas. In turn, bananas have far more to do with food, fruit, diet, health, and so on. Working at broadband speeds, such a system could grow in capacity and competence at an alarming rate. All that's needed is a source of concepts to add. But how do we use the resultant wisdom to create strong AI? Cilibrasi and Vitanyi have already found three uses for their discovery.


The first involves using a program called Complearn to find clusters of closely related concepts. When presented with the names of colours and numbers, Complearn organises them so terms that are definitely colours appear on

one side of its output diagram and numbers occupy the other. More general terms such as 'zero', 'white' and 'black' appear nearer the middle, and terms such as 'light', 'dark', 'big' and 'small' occupy the centre.

Incredibly, when given a list of paintings by 17th century Dutch masters, Complearn correctly grouped them by artist, without being told explicitly who had painted which work.

The researchers have used a machine-learning algorithm called a Support Vector Machine (SVM) to deduce new knowledge from the facts stored. An SVM acts like a neural network; it looks for patterns in data and when it is presented with data it has not seen before, it can find matches in patterns to arrive at a result. Cilibrasi and Vitanyi have used this technique to establish what prime numbers are. They say their system can 'get' what prime numbers are without being told explicitly. It can also work out when a situation is an emergency, when something is likely to be an electrical engineering term, and when it is likely to be religious.

In a third application, the system tries to translate between English and Spanish. The knowledge base learns to associate words in both languages, and even their grammatical quirks, to translate new tracts of text.



Cilibrasi and Vitanyi think that by extracting meaning from Google a complex enough system could operate usefully just by being given examples of whatever you wish it to learn – exactly as a human does. The researchers point out other potential examples, including automatic language acquisition, whereby software learns a new language on its own. It might do this by example, or by listening to someone speak.

Personal services

Artificial intelligence might conjure up images of robot pets or talking cars, but there are plenty more practical applications in store. Suppose you want to book a family holiday in the Lake District. The logistics of finding out what accommodation is available and when, booking time off work, arranging transport, planning trips out and things to do are all time-consuming, not to mention stressful.

In future, you might simply tell a personal software agent to book the holiday, stipulating any restrictions such as school holidays. Its first task might be to contact your children's school. Knowing that you need to go during the school holidays, the agent would take a list of dates to check accommodation availability.

Next, by looking up the destination area in a geographical database, it could figure

out which hotels, guesthouses or holiday cottages with websites fit your requirements, because it knows what constitutes holiday accommodation. Once located, it could access the various online booking systems to obtain a list of free dates and prices.

The next stop might be your company's human resources system, where your personal software agent would use the time-booking system to check your holiday entitlement. Since driving to and from a holiday can be tiresome, you might wish to travel by bus or train and hire a car locally. If no cars are available, the agent might study a local transport database to find an appropriate substitute, such as a taxi service.

What the agent might potentially bring you back is a holiday itinerary, saving you the time and stress of trying to coordinate everything manually.

INDEPENDENCE DAY

It seems Lenat was right in his claim that a machine must have access to all the common sense we take for granted to stand a chance of being thought of as intelligent. Using Cyc, he also showed that the more complex a view of the world a program has, the more intelligent its behaviour becomes.

Now that software has access to the ever-growing web, it's tempting to speculate that it may be only a matter of time before complex, human-style reasoning is simulated using computers. But surely all it will be doing is obeying the set of common-sense rules that we follow? On the surface, such a system wouldn't have the capacity for creative thoughts, for instance.

Still, occasional slight mutations to concepts and relationships within its growing knowledge base might be all it takes for new and independent thought to start emerging. Eventually, we wouldn't be able to say with any certainty how such a system arrived at a solution or an opinion about something. Maybe then we might start to wonder if common sense is all that's needed for intelligent thought after all. **CS**